

Interventions and Causal Inference

Frederick Eberhardt¹ and Richard Scheines
Department of Philosophy
Carnegie Mellon University

Abstract

The literature on causal discovery has focused on interventions that involve randomly assigning values to a single variable. But such a randomized intervention is not the only possibility, nor is it always optimal. In some cases it is impossible or it would be unethical to perform such an intervention. We provide an account of “hard” and “soft” interventions, and discuss what they can contribute to causal discovery. We also describe how the choice of the optimal intervention(s) depends heavily on the particular experimental set-up and the assumptions that can be made.

Introduction

Interventions have taken a prominent role in recent philosophical literature on causation, in particular in work by James Woodward (2003), Christopher Hitchcock (2005), Nancy Cartwright (2006, 2002) and Dan Hausman and James Woodward (1999, 2004). Their work builds on a graphical representation of causal systems developed by computer scientists, philosophers and statisticians called “Causal Bayes Nets” (Pearl, 2000; Spirtes, Glymour, Scheines (hereafter SGS), 2000). The framework makes interventions explicit, and introduces two assumptions to connect qualitative causal structure to sets of probability distributions: the Causal Markov and Faithfulness assumptions. In his recent book, *Making Things Happen* (2003), Woodward attempts to build a full theory of causation on top of a theory of interventions. In Woodward’s theory, roughly, one variable X is a *direct cause* of another variable Y if there exists an intervention on X such that if all other variables are held fixed at some value, X and Y are associated. Such an account assumes a lot about the sort of intervention needed, however, and Woodward goes to great lengths to make the idea clear. For example, the intervention must make its target independent of its other causes, and it must directly influence only its target, both of which are ideas difficult to make clear without resorting to the notion of direct causation.

Statisticians have long relied on intervention to ground causal inference. In *The Design of Experiments* (1935), Sir Ronald Fisher considers one treatment variable (the purported cause) and one or more effect variables (the purported effects). This approach has since been extended to include multiple treatment and effect variables in experimental designs such as Latin and Graeco-Latin squares and Factor experiments. In all such cases, however, one must designate certain variables as potential causes (the treatment variables) and others as potential effects (the outcome variables), and inference begins with a randomized assignment (intervention) of the potential cause. Similarly, the

¹ The first author is supported by the Causal Learning Collaborative Initiative supported by the James S. McDonnell Foundation. Many aspects of this paper were inspired by discussions with members of the collaborative.

framework for causal discovery developed by Donald Rubin (1974, 1977, and 1978) assumes there is a treatment variable and that inferences are based on samples from randomized trials.

Although randomized trials have become de facto the gold standard for causal discovery in the natural and behavioral sciences, without such an a priori designation of causes and effects Fisher's theory is far from complete. First, without knowing ahead of time which variables are potential causes and which effects, more than one experiment is required to identify the causal structure – but we have no account of an optimal sequence of experiments. Second, we don't know how to statistically combine the results of two experiments involving interventions on different sets of variables. Third, randomized assignments of treatment are but one kind of intervention. Others might be more powerful epistemologically, or cheaper to execute, or less invasive ethically, etc.

The work we present here describes two sorts of interventions (“structural” and “parametric”) that seem crucial to causal discovery. These two types of interventions form opposite ends of a whole continuum of ‘harder’ to ‘softer’ interventions. The distinction lines up with interventions of different “dependency”, as presented by Korb (2004). We then investigate the epistemological power of each type of intervention without assuming that we can designate ahead of time the set of potential causes and effects. We give results about what can and cannot be learned about the causal structure of the world from these kinds of interventions, and how many experiments it takes to do so.


Causal Discovery with Interventions

Causal discovery using interventions not only depends on what kind of interventions one can use, but also on what kind of assumptions one can make about the models considered. We assume that the causal Markov and faithfulness assumptions are satisfied (see Cartwright (2001, 2002) and Sober (2001) for exceptions and Steel (2005) and Hoover (2003) for sample responses). The causal Markov condition amounts to assuming that the probability distribution over the variables in a causal graph factors according to the graph as so:

$$P(X) = \prod_{X_i \in X} P(X_i \mid \text{parents}(X_i))$$

where the parents of X_i are the immediate causes of X_i , and if there are no parents then the marginal distribution $P(X_i)$ is used. For example,

Causal Graph



Factored Joint Distribution

$$P(T, A, HO) = P(A) P(T | A) P(HO | A)$$

The faithfulness assumption says that the causal graph represents the causal independence relations true in the population, i.e. that there are no two causal paths that cancel each other out. However, several other conditions have to be specified in order to formulate a causal discovery problem precisely.

It is often assumed that the causal models under consideration satisfy causal sufficiency and acyclicity. Causal sufficiency is the assumption that there are no unmeasured common causes of any pair of variables that are under consideration (no latent confounders). Assuming causal sufficiency is unrealistic in most cases, since we rarely measure all common causes of all of our pairs of variables. However, the assumption has a large effect on a discovery procedure, since it reduces the space of models under consideration substantially. Assuming acyclicity prohibits models with feedback. While this may also be an unrealistic assumption – in many situations in nature there are feedback cycles – this is beyond the scope of this paper.

With these assumptions in place we turn to “structural” and “parametric” interventions. The typical intervention in medical trials is a “structural” randomization of one variable. Subjects are *assigned* to the treatment or control group based on a random procedure, e.g. a flip of a coin. We call the intervention “structural” because it alone completely determines the probability distribution of the target variable. It makes the intervened variable independent of its other causes and therefore changes the causal structure of a system before and after such an intervention. In a drug trial, for example, a fair coin determines whether each subject will be assigned one drug or another or a placebo. The assignment of treatment is independent of any other factor that might cause the outcome.² Randomization of this kind ensures that, at least in the probability distribution true of the population, such a situation does not arise.

In Causal Bayes Nets a *structural* intervention is represented as an exogenous variable I (a variable without causes) with two states (*on/off*) and a single arrow into the variable it manipulates.³ When I is *off*, the passive observational distribution obtains over the variables. When I is *on*, all other arrows incident on the intervened variable are removed, and the probability distribution over the intervened variable is a determinate function of


² For example, treatment and age should not be associated, as age is likely a cause of response to almost any drug.

³ For more detail, see SGS (2000), chapter 7, see also appendix here.


the intervention only. This property underlies the terminology “structural.”⁴ If there are multiple simultaneous structural interventions on variables in the graph, the manipulated distribution for each intervened variable is independent of every other manipulated distribution⁵, and the edge breaking process is applied separately to each variable. This implies that *all* edges between variables that are subject to an intervention are removed. After removing all edges from the original graph incident to variables that are the target of a structural intervention, the resulting graph is called the *post-manipulation graph* and represents what is called the manipulated distribution over the variables.

More formally, we have the following definition for a *structural* intervention I_s on a variable X in a system of variables V :

- I_s is a variable with two states (*on/off*).
- When I_s is *off*, the passive observational distribution over V obtains.
- I_s is a direct cause of X and only X .
- I_s is exogenous⁶, i.e. uncaused.
- When I_s is *on*, I_s makes X independent of its causes in V (breaks the edges that are incident on X) and *determines* the distribution of X ; that is, in the factored joint distribution $P(V)$, the term $P(X \mid \text{parents}(X))$ is replaced with the term $P(X \mid I_s)$, all other terms in the factorized joint distribution are unchanged.



$$P(T, A, HO \mid I=off) = P(A) P(T \mid A) P(HO \mid A, T)$$



$$P(T, A, HO \mid I=on) = P(A) P(T \mid I=on) P(HO \mid A, T)$$

Causal Bayes Net representing a “**structural**” intervention (e.g. a medical trial) where Treatment (T) and Health Outcome (HO) are confounded by an unmeasured common cause Age (A). When the intervention variable I is set to *off*, the passive observational distribution obtains over the variables. When it is set to *on*, the structural intervention breaks any arrows incident on the treatment variable, thereby destroying any correlation due to unmeasured common causes.

⁴ Pearl refers to these interventions as “surgical” for the same reason. They are also sometimes referred to as “hard” or “ideal” interventions. Korb (2004) refers to them as “independent” interventions.

⁵ This is an assumption we make here for the theorems that follow. It is not an assumption that is necessary in general. There may well be so-called “fat-hand” interventions that imply correlated interventions on two variables. But in such cases the discovery procedure has to be adapted and the following results do not hold in general.

⁶ We do not take ‘exogenous’ to be synonymous with ‘uncaused’ in general, but rather to represent a weaker notion. However, for brevity we will skip that discussion here.

The epistemological advantages of structural interventions on one variable are at least the following:

- No correlation between the manipulated variable and any other non-manipulated variable in the resulting distribution is due to an unmeasured common cause (confounder).
- The structural intervention provides an ordering that allows us to distinguish the direction of causation, i.e. it distinguishes between $A \rightarrow B$ and $A \leftarrow B$.
- The structural intervention provides a fixed known distribution over the treatment variable that can be used for further statistical analysis, such as the estimation of a strength parameter of the causal link.



This is not the only type of intervention that is possible or informative, however. There may also be “soft” interventions that do not remove edges, but simply modify the conditional probability distributions of the intervened upon variable. In a causal Bayes net, such an intervention would still be represented by an exogenous variable with a single arrow into the variable it intervenes on. Again, when it is set to *off*, the passive observational distribution obtains; but when it is set to *on*, the distribution of the variable conditional on its causes (graphical parents) is changed, but their causal influence (the incoming arrows) are not broken. We refer to such an intervention as a “parametric” intervention, since it only influences the parameterization of the conditional probability distribution of the intervened variables on its parents, while it still leaves the causal structure intact.⁷ The conditional distribution of the variable still remains a function of the variable's causes (parents).

More formally, we have the following definition for a *parametric* intervention I_p on a variable X in a system of variables V :

- I_p is a variable with two states (*on/off*).
- When I_p is *off*, the passive observational distribution over V obtains.
- I_p is a direct cause of X and only X .
- I_p is exogenous, i.e. uncaused.
- When I_p is *on*, I_p does not make X independent of its causes in V (does not break the edges that are incident on X). In the factored joint distribution $P(V)$, the term $P(X \mid \text{parents}(X))$ is replaced with the term $P(X \mid \text{parents}(X), I_p = \text{on})$,⁸ and otherwise all terms are unchanged.

⁷ As mentioned above, Korb (2004) refers to such an intervention as a “dependent” intervention. It has also been referred to as “soft” (Campbell, 2006) or “conditional”.

⁸ Obviously, $P(X \mid \text{parents}(X)) \neq P(X \mid \text{parents}(X), I_p = \text{on})$.



Causal Bayes Net representing a “**parametric**” intervention. When the intervention variable I is set to *off*, the passive observational distribution obtains over the variables. When it is set to *on*, the parametric intervention does not break any arrows incident on the intervened variable, but only changes the conditional distribution of that variable. So, $P(IC|SES) \neq P(IC|SES, I=on)$. In contrast to structural interventions, parametric interventions do not break correlations due to unmeasured common causes.

There are several ways to instantiate such a parametric intervention. If the intervened variable is a linear (or additive) function of its parents, then the intervention could be an additional linear factor. For example, if the target is *income*, the intervention could be to boost the subject’s existing income by \$10,000/year. In the case of binary variables, the situation is a little more complicated, since the parameterization over the other parents must be changed, but even here it is possible to perform a parametric intervention, e.g. by inverting the conditional probabilities of the intervened variable when the parametric intervention is switched to *on*.⁹

In practice parametric interventions can arise for at least two reasons. In some cases it might be impossible to perform a structural intervention. Campbell (2006) argues that one cannot perform a structural intervention on mental states, since it is impossible to make a mental state independent of its other causes: One can make a subject believe something, but that belief will not be independent of, for example, prior beliefs and experiences. Or if it is, then it is questionable whether one can still speak of the *same* subject that has this belief. Also, parametric interventions may arise when the cost of a structural intervention is too high or when a structural intervention of a particular variable would be unethical. For example, it is in principle possible to randomly assign (structural intervention) income, but the cost would be enormous and it is unethical to assign someone an income that would be insufficient for survival, while denying the participants their normal income. In such cases one might instead add a fixed amount to the income of the subjects (parametric intervention). Income would then remain a function of its prior causes (e.g. education, parental income, socio economic status (SES) etc.), but would be modified by the parametric intervention.


⁹ We are grateful to Jiji Zhang for this example.

Naturally, experiments involving parametric interventions provide different information about the causal structure than those that use structural ones. In particular, they do not destroy any of the original causal structure. This has the advantage that causal structure that would be destroyed in structural interventions might be detectable, but it has the disadvantage that an association due to a (potentially unmeasured) common cause is not broken. In the following we provide an account of the implications these different types of interventions have on what we can learn about the causal structure, and how fast we can hope to do so.

Results

While a structural intervention is extremely useful to test for a direct causal link between two variables (this is the focus in the statistics literature), it is not straight forwardly the case that structural interventions on single variables provide an efficient strategy for discovering the causal structure among several variables. The advantage it provides, namely making the intervened upon variable independent of its other causes, is also its drawback. In general, we still want a theory of causal discovery that does not rely upon an a priori separation of the variables into treatment and effect as is assumed in statistics. Even time ordering does not always imply information about such a separation, since we might only have delayed measures of the causes.

Faced with a setting in which any variable may be a cause of any other variable, a structural intervention of the wrong variable might then not be informative about the true causal structure, since even the manipulated distribution could have been generated by several different causal structures.



For example, consider the above Figure. Suppose the true but unknown causal graph is (1). A structural intervention on C would make the pairs A-C and B-C independent, since the incoming arrows on C are broken in the post-manipulation graph (2). The problem is that the information about the causal influence of A and B on C is lost. Note also, that an association between A and B is detected but the direction of the causal influence cannot be determined (hence the representation by an undirected edge). The manipulated distribution could as well have been generated by graph (3), where the true causal graph has no causal links between A and C or B and C. Hence, structural interventions also create Markov equivalence classes of graphs, that is, graphs that have a different causal structure, but imply the same conditional independence relations. (1) and (3) form part of an interventional Markov equivalence class under a structural

intervention on C (they are not the only two graphs in that class, since the arrow between A and B could be reversed as well). Discovering the true causal structure using structural interventions on a single variable, and to be guaranteed to do so, requires a sequence of experiments to partition the space of graphs into Markov equivalence classes of unique graphs. Note that a further structural intervention on A in a second experiment would distinguish (1) from (3), since A and C would be correlated in (1) while they would be independent in (3).


Eberhardt, Glymour, and Scheines (2006) showed that for N causally sufficient variables N-1 experiments are sufficient and in the worst case necessary to discover the causal structure among a causally sufficient set of N variables if at most one variable can be subjected to a structural intervention per experiment assuming faithfulness. If multiple variables can be randomized simultaneously and independently in one experiment, this bound can be reduced to $\log(N) + 1$ experiments (Eberhardt et al, 2005). These bounds both assume that an experiment specifies a subset of the variables under consideration that are subject to an intervention and that each experiment returns the independence relations true in the manipulated population, i.e. issues of sample variability are not addressed.

Parametric interventions do not destroy any of the causal structure. However, if only a single parametric intervention is allowed, then there is no difference in the number of experiments between structural and parametric interventions:

*Theorem 1: N-1 experiments are sufficient and in the worst case necessary to discover the causal relations among a causally sufficient set of N variables if only one variable can be subject to a **parametric** intervention per experiment. (Proof sketch in Appendix)*

For experiments that can include simultaneous interventions on several variables, however, we can decrease the number of experiments from $\log(N) + 1$ to a *single* experiment when using parametric interventions:

*Theorem 2: One experiment is sufficient and (of course) necessary to discover the causal relations among a causally sufficient set of N variables if multiple variables can be simultaneously and independently subjected to a **parametric** intervention per experiment. (Proof sketch in Appendix)*



The following example, illustrated in the above figure, explains the result: The true unknown complete graph among the variables A, B and C is shown on the left. In one experiment, the researcher performs simultaneously and independently a parametric

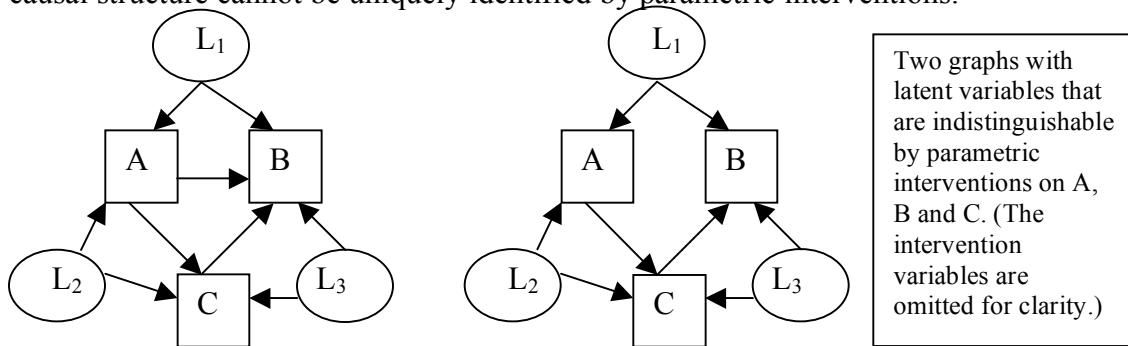
intervention on A and B (I_A and I_B , respectively, shown on the right). Since the interventions do not break any edges, the graph on the right represents the post-manipulation graph. Note that A, B and I_B form an unshielded collider,¹⁰ as do C, B and I_B . These can be identified¹¹ and hence determine the edges and their directions A to B and C to B. The edge A to C can be determined since (i) A and C are dependent for all possible conditioning sets, but (ii) I_A , A and C do *not* form an unshielded collider. Hence we can conclude that (from (i)) there must be an edge between A and C and (from (ii)) that it must be directed away from A. We have thereby managed to discover the true causal graph in one experiment. Essentially, adjacencies can be determined from observational data alone. The parametric interventions set up a “collider test” for each triple I_X , X and Y with X – Y adjacent, which orients the X – Y adjacency.

Discussion

These results indicate that the advantage of parametric interventions lies with the fact that they do not destroy any causal connections.

Number of experiments for different types of interventions	Single Intervention per experiment	Multiple simultaneous interventions per experiment
Parametric Interventions	$N - 1$	1
Structural Interventions	$N - 1$	$\log_2(N) + 1$

The theorems tempt the conclusion that parametric interventions are always better than structural interventions. But this would be a mistake since the theorems hide the cost of this procedure. First, determining the causal structure from parametric interventions requires more conditional independence tests with larger conditioning sets. This implies that more samples are needed to obtain a similar statistical power on the independence tests as in the structural intervention case. Second, the above theorems only hold in general for causally sufficient sets of variables. A key advantage of randomized trials (structural interventions) is their robustness against latent confounders (common causes). Parametric interventions are not robust in this way, since they do not make the intervened variable independent of its other causes. This implies that there are cases for which the causal structure cannot be uniquely identified by parametric interventions.




¹⁰ Variables X, Y and Z form an unshielded collider if X is a direct cause of Y, Z is a direct cause of Y and X is not a direct cause of Z and Z is not a direct cause of X. Unshielded colliders can be identified in statistical data, since X and Z are unconditionally independent, while they are dependent conditional on Y.

¹¹ See previous footnote.

The two graphs above over the observed variables A, B and C with latent common causes L_1 , L_2 and L_3 are indistinguishable given parametric interventions on A, B and C. There is no conditional independence relation between the variables (including the intervention nodes omitted for clarity in the figure) that distinguishes the two graphs¹².

However, it is not always the case that causal insufficiency renders parametric interventions useless as the following example shows.



Even in the presence of a latent variable, this causal structure can be uniquely determined with parametric interventions on X and Y in a single experiment.

The parametric intervention I_X on X will not break the association between X and Y that is due to the unmeasured common cause L. This does not mean that the edge X to Y cannot be identified, since (a) if there were no edge between X and Y, then I_X and Y *would not* be associated, and (b) if the edge were from Y to X, then I_Y and X *would* be associated.¹³ As the two examples show, parametric interventions are *sometimes* subject to failure when causal sufficiency cannot be assumed, but this depends very much on the complexity of the model.¹⁴

It follows that the assumption of causal sufficiency does not compensate entirely for what parametric interventions lack in robustness to unmeasured common causes. First, the results for structural interventions also depend on causal sufficiency.¹⁵ A logarithmic bound on the number of experiments when using structural interventions would not be achievable in the worst case if there are latent common causes. Second, there are cases (as the example above shows) where parametric interventions are still sufficient to recover the causal structure even if there are latent common causes.¹⁶ In some cases one can even identify the location of the latent variables. However, this is not always the case

¹² In order to prove this claim, a huge number of conditioning sets have to be checked, which we did using the Tetrad program (causality lab: <http://www.phil.cmu.edu/projects/causality-lab/>), but omit here for brevity.

¹³ Note, that it is not *simply* the creation of unshielded colliders that is doing the work here. If there were no parametric intervention on Y, we could not distinguish between $X \leftarrow Y$ and no edge between X and Y at all, *even though* the first case would create an unshielded collider $I_X \rightarrow X \leftarrow Y$. In the presence of latent variables, creating colliders is necessary, but not sufficient (as the more complex example shows).

¹⁴ Note also the close similarity between parametric interventions and instrumental variables as they are used in economics. However, in economics – as in statistics – there is the general assumption that some kind of order is known between the variables, so that one only tests for presence of an adjacency, while its direction is determined by background knowledge.

¹⁵ See Eberhardt et al. (2005).

¹⁶ It is not the simplicity of the second example (with just two variables) that allows for the identifiability of the structure when using parametric interventions in the presence of latent common causes. In fact, if we remove the edge $C \rightarrow B$ (instead of $A \rightarrow B$) in the three variable example above, the graph would be distinguishable from the complete graph among A, B and C even if there are latent common causes for every pair of variables. However, it may not form a singleton Markov equivalence class under parametric interventions.

and we do not have any general theorem of these cases to report at this stage. What seems to play a much larger role than causal sufficiency is the requirement of exogeneity of the intervention. It results in additional independence constraints (creating unshielded colliders) whose presence or absence make associations identifiable with direct causal connections.

The moral is that while the types of interventions are not independent of the assumptions that are made about causal sufficiency, they are not interchangeable with them either. It is not the case that causal sufficiency and parametric interventions achieve the same search capability as causal insufficiency and structural interventions. The trade-offs are rather subtle and we have only shown a few examples of the interplay of some of the assumptions. But there does not seem to be any general sense in which one can speak of “weaker” or “stronger” interventions with regard to their epistemological power.

Appendix:

Intervention Variables

We represent interventions in causal Bayes nets by so-called policy variables, which have two states. These policy variables are not causal variables in the sense of the other variables in the causal graph (although they may model instrumental variables): they need not have a marginal probability distribution over their values. By “switching” the policy variable to one of its states we refer to the decision to perform an intervention. In the case of a parametric intervention, the policy variable creates unshielded colliders, while we claim that this is not the case for structural interventions. This will need some clarification, since this aspect is very relevant to the discovery procedures.

If our data contains samples both for when a variable X is subject to a structural intervention *and* when it is passively observed, then – using the sample distribution of I_s – we will find that even in the case of structural interventions, $I_s \rightarrow X \leftarrow Y$ forms an unshielded collider. The adjacency is obtained from the subsample where $I_s = \text{off}$, while the direction is determined when $I_s = \text{on}$. This sample constitutes a mixture of populations, one manipulated and one unmanipulated, as we may find it in a randomized trial with a control group that is not subject to any intervention. However, in studies where each condition in the randomized trial involves an intervention¹⁷ (e.g. comparative medical studies) we do not have such a passively observed control group that would capture the unmanipulated structure. So the interventions are different (with regard to the manipulated distribution they impose) in different conditions, but if both conditions constitute *structural* interventions, then the causal structure cannot be recovered from the sample in all cases. Furthermore, if we perform multiple simultaneous structural interventions and only obtain data where either *all* intervention variables are *on* or all intervention variables are *off* for each sample, then again we cannot recover the causal structure in all cases.

¹⁷ In this case the “on”-state ($I_s = \text{on}$) of the intervention variable would have to be augmented to different on-states.

In contrast, we can consider the unshielded collider in experimental designs involving *parametric* interventions since we can discover the unshielded collider even in a data set where the sample $I_p = on$ for all samples, as long as there are different “on”-states. And we can discover unshielded colliders in the case of multiple simultaneous parametric interventions, i.e. when we only obtain data where either *all* intervention variables are *on* or all intervention variables are *off* for each sample. The key is that parametric interventions can be combined independently and performed simultaneously without interfering, while this is not the case for structural interventions.

Theorem 1: (Proof Sketch)

Let each of the $N-1$ experiments E_i , with $0 < i < N$ consist of a parametric intervention on X_i . In each case, the intervention variable I_i forms an unshielded collider with any cause of X_i . Hence, for any variable Y , where Y and I_i are unconditionally independent, but dependent conditional on C union $\{Y\}$ for all conditioning sets C , we know that Y is a cause of X_i . Further, in each experiment we check whether X_i and X_N (which is not subject to an intervention) are dependent for all conditioning sets. If so, then X_i is a cause of X_N . Since we perform a parametric intervention on $N-1$ variables, all causes of these $N-1$ variables can be discovered, and since we check for each variable whether it is a cause of X_N , all its causes are determined as well. Hence, $N-1$ experiments are sufficient to discover the causal structure.

$N-1$ experiments are in the worst case necessary, since $N-2$ parametric interventions would imply that two variables are not subject to an intervention. This would make it (in the worst case) impossible to determine the direction of any edge between them, if there were one.

Theorem 2: (Proof Sketch)

Since the parametric interventions described in the proof of the previous theorem do not interfere with each other, they can be performed all at the same time. That is, in a single experiment $N-1$ variables would be subject to a parametric intervention and the causal structure could be discovered all in one go. One experiment is necessary, since -- in the case of a complete graph -- there would not be any unshielded colliders that would allow for the determination of the direction of any causal link between the variables. This is where the parametric $N-1$ interventions are necessary. The theorem can also be derived simply from a theorem on rigid indistinguishability in SGS¹⁸.

Bibliography:

Campbell, J. (2006) An Interventionist Approach to Causation in Psychology, in Alison Gopnik and Laura Schulz (eds.), *Causal Learning: Psychology, Philosophy and Computation*, Oxford: Oxford University Press, in press.

¹⁸ Theorem 4.6, Chapter 4.

- Cartwright, N. (1983). *How the Laws of Physics Lie*, Oxford University Press
- Cartwright, N. (2001). What is Wrong with Bayes Nets?, in *The Monist*, pp 242-264
- Cartwright, N. (2002). Against Modularity, the Causal Markov Condition and Any Link Between the Two, in *British Journal of Philosophy of Science*, 53: 411-453
- Cartwright, N. (2006). From Metaphysics to Method: Comments on Manipulability and the Causal Markov Condition, in *British Journal of Philosophy of Science*, 57: 197-218
- Chickering. D. M. (2002). Learning Equivalence Classes of Bayesian-Network Structures. *Journal of Machine Learning Research*, 2:445-498.
- Eberhardt, F., Glymour, C., Scheines. R., (2005). On the Number of Experiments Sufficient and in the Worst Case Necessary to Identify All Causal Relations Among N Variables, in *Proceedings of the 21 st Conference on Uncertainty and Artificial Intelligence*, Fahiem Bacchus and Tommi Jaakkola (editors), AUAI Press, Corvallis, Oregon, pp. 178-184.
- Eberhardt, F., Glymour, C., Scheines. R., (2006). N-1 Experiments Suffice to Determine the Causal Relations Among N Variables, in *Innovations in Machine Learning*, Holmes, Dawn E.; Jain, Lakhmi C. (Eds.), Theory and Applications Series: Studies in Fuzziness and Soft Computing, Vol. 194, Springer-Verlag.
- Fisher, R. (1935). *The Design of Experiments*. Hafner.
- Glymour, C. & Cooper. G (1999). Computation, Causation, and Discovery, in AAAI Press and MIT Press, Cambridge, MA
- Hausman, D. and Woodward, J. (1999). Independence, Invariance, and the Causal Markov Condition, in *British Journal for the Philosophy of Science*, 50: pp. 521–83.
- Hausman, D. and Woodward, J. (2004). Modularity and the Causal Markov Condition: A Restatement, in *British Journal for the Philosophy of Science*, 55: pp. 147–61.
- Hitchcock C. (2006), On the Importance of Causal Taxonomy in A. Gopnik and L. Schulz (eds.), *Causal Learning: Psychology, Philosophy and Computation* (New York: Oxford University Press).
- Hoover, K. (2003) Nonstationary Time Series, Cointegration, and the Principle of Common Cause, in *British Journal of Philosophy of Science*, 54, pp 527-551
- Hume, D. (1748). *An Enquiry Concerning Human Understanding*. Oxford, England: Clarendon

- Kant, I. (1781) Critique of Pure Reason
- Korb, K., Hope, L., Nicholson, A. & Axnick, K. (2004), Varieties of Causal Intervention, in *Proceedings of the Pacific Rim International Conference on AI*, Springer
- Lewis, D. (1973) Counterfactuals. Harvard University Press
- Mackie, J. (1974) The Cement of the Universe, Oxford University Press, New York.
- Meek, C. (1996). PhD Thesis, Department of Philosophy, Carnegie Mellon University
- Mill, J.S. (1843, 1950). Philosophy of Scientific Method. New York: Hafner.
- Murphy, K.P. (2001). Active Learning of Causal Bayes Net Structure, Technical Report, Department of Computer Science, U.C. Berkeley.
- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems*, San Mateo, CA. Morgan Kaufmann.
- Pearl, J. (2000). *Causality*, Oxford University Press.
- Rubin, D (1974). Estimating causal effects of treatments in randomized and non-randomized studies, in *Journal of Educational Psychology*, 66: 688-701
- Rubin, D (1977). Assignment to treatment group on the basis of covariate, in *Journal of Educational Statistics*, 2: 1-26
- Rubin, D (1978). Bayesian Inference for causal effects: The role of randomizations, in *Annals of Statistics*, 6: 34-58
- Salmon, W. (1984) Scientific Explanation and the Causal Structure of the World. Princeton University Press
- Sober, E. (2001), Venetian Sea Levels, British Bread Prices, and the Principle of Common Cause, *British Journal of Philosophy of Science*, 52, pp. 331-346
- P. Spirtes, C. Glymour and R. Scheines, (1993). *Causation, Prediction and Search*, Springer Lecture Notes in Statistics, 1993; 2nd edition, MIT Press, 2000.
- P. Spirtes, C. Meek & T. Richardson, (1999). An Algorithm for Causal Inference in the Presence of Latent Variables and Selection Bias. In *Computation, causation and discovery*. MIT Press.
- P. Spirtes, C. Glymour, R. Scheines, S. Kauffman, W. Aimalie, and F. Wimberly, (2001). Constructing Bayesian Network Models of Gene Expression Networks from

Microarray Data, in *Proceedings of the Atlantic Symposium on Computational Biology, Genome Information Systems and Technology*, Duke University, March.

Suppes, P. (1970), *A Probabilistic Theory of Causality*. North-Holland, Amsterdam.

Steel, D. (2005), Indeterminism and the Causal Markov Condition, in *British Journal of Philosophy of Science*, 56: 3-26

S. Tong and D. Koller, (2001). Active Learning for Structure in Bayesian Networks, *Proceedings of the International Joint Conference on Artificial Intelligence*.

J. Woodward, (2003). *Making Things Happen*, Oxford University Press